REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED *(From - To)*
02-02-2001		N/A

4. TITLE AND SUBTITLE
Export Control of High Performance Computing: Analysis and Alternative Strategies

5a. CONTRACT NUMBER
N/A

5b. GRANT NUMBER
N/A

5c. PROGRAM ELEMENT NUMBER
N/A

6. AUTHOR(S)
Dr. Charles J. Holland
Dr. Delores M. Etter
Mr. John Grosh

5d. PROJECT NUMBER
N/A

5e. TASK NUMBER
N/A

5f. WORK UNIT NUMBER
N/A

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

Office of the Deputy Under Secretary of Defense
(Science and Technology}
3040 Defense Pentagon
Washington, DC 20301

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Same as above.

10. SPONSOR/MONITOR'S ACRONYM(S)
ODUSD(S&T)

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION / AVAILABILITY STATEMENT

Distribution Statement A

13. SUPPLEMENTARY NOTES

14. ABSTRACT

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			James McDonald
Unclassified	Unclassified	Unclassified	None	9 incl covers	19b. TELEPHONE NUMBER *(include area code)* (703) 697-8535

DEFENSE SCIENCE AND TECHNOLOGY TECHNICAL REPORTS

Export Control of High Performance Computing:
Analysis and Alternative Strategies

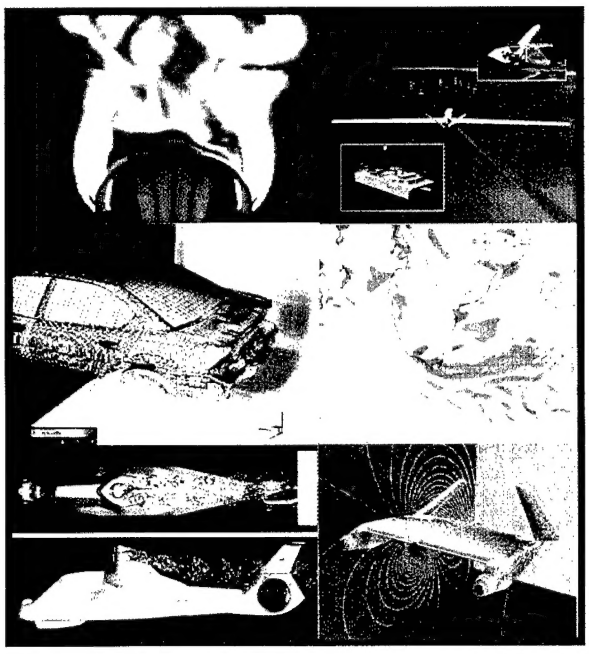

February 2, 2001
Cleared for Public Release

Office of the Deputy Under Secretary of Defense (Science and Technology)
3080 Defense Pentagon

Report on

Export Control of High Performance Computing:

Analysis and Alternative Strategies

Dr. Delores M. Etter
Deputy Under Secretary of Defense (Science and Technology)

Dr. Charles J. Holland
Director, Information Technology, Office of the Deputy
Under Secretary of Defense (Science and Technology)

Mr. John Grosh
Senior Staff Specialist for High Performance Computing and Software,
Office of the Deputy Under Secretary of Defense (Science and Technology)

3080 Defense, Pentagon
Washington, DC 20301-3080

February 2, 2001

Executive Summary

High performance computing has historically played an important role in the ability of the United States to develop and deploy a wide range of national security capabilities, such as stealth aircraft, sonar arrays, and high-energy rocket fuels. Therefore, it is critical that the United States stay ahead technically in this important area. Export controls on computer hardware are one of several strategies used to ensure U.S. superiority in high performance computing. However, rapid advances in computer technology have limited the government's ability to prevent the export of high performance computing hardware to potential adversaries. Furthermore, controls that once restricted the export of high-end supercomputers now restrict the export of low- and mid-range servers. Denial of access to growing third-world markets, burgeoning foreign computer manufacturing capability, and increasing foreign demand for computers give rise to industry concerns over the potential loss of American dominance in the world computer market. These issues are seen as a threat to both U.S. economic security and national defense [1].

In light of these concerns, the Deputy Under Secretary of Defense (Science and Technology) was asked to conduct a study to develop alternative export control strategies for high performance computing. The goal of this study was to develop strategies that protect national security interests in high performance computing while not endangering U.S. dominance of the computer industry. This report summarizes the findings of this work, completed in November 2000. The current hardware control strategy is based on a measure of computer performance known as MTOPS, millions of theoretical operations per second. This metric, implemented in 1991, served as an effective basis for export controls until the late 1990s [2]. Our analysis shows that by late 1999, MTOPS controls could be easily circumvented and were no longer effective. Thus, our recommended strategy is to abandon MTOPS controls for high performance computers. After evaluating a number of alternative hardware control strategies, we concluded that it is no longer feasible to control high performance computer hardware. Instead, we recommend a focus on protecting our software used specifically for national security applications.

High performance computing involves more that just hardware. Highly sophisticated and specialized applications software is also required to effectively utilize high performance computer systems. Military and dual-use applications software represents an important national security investment. Such software is acquired through extensive software development, testing, and validation with experiments (e.g., expensive live fire tests) and provides the defense community with a capability that is not easily duplicated by an adversary. With the loss of effective hardware controls, we conclude that emphasis must shift to software protection and controls as the principal means of restricting foreign exploitation of high performance computing. Our proposed strategy includes research and development in technologies that prevent unauthorized use of applications software, as well as training and information dissemination to increase the effectiveness of existing software control policies.

Table of Contents

List of Tables and Figures

I. Export Control of High Performance Computing Hardware

A. Background

Government regulations for export control of computer systems are currently based on a measure of computer systems performance, known as composite theoretical performance, and a threat assessment of the destination country. Composite theoretical performance is measured in units of millions of theoretical operations per second (MTOPS) and is dependent on a system's processor speed and word length, number of processors, and interconnect configuration [3]. Threat assessments are categorized into tiers, as listed in Table 1, based upon the potential threat posed by these countries to the interests of the United States.

A license exception is an authorization for exporting or re-exporting items that would otherwise require a license under the prohibitions of the Export Administration Regulations. For the export of computers, license exception limits are based on MTOPS. Exception limits for the past nine years are described in Table 2. Licenses are required by the vendors to export computers rated above these limits. For systems below the exception limits and above 6,500 MTOPS, vendors are granted an exception to the licensing requirement, but are required to keep records of exports and report these to the government every six months. For Tier 3 exports, the export of computers between 12,500[1] and 85,000 MTOPS are subject to a ten-day government review to obtain a license exception. There is a virtual embargo on exports to Tier 4 countries.

Tier Level	Threat Description	Countries
Tier 1	Allies and friendly nations	Western Europe, Japan, Canada, Mexico, Australia, New Zealand, Hungary, Poland, the Czech Republic, and Brazil
Tier 2[2]	Medium-risk destinations	South and Central America, South Korea, Association of South Asian States, Slovenia, most of Africa, Romania
Tier 3	High-risk destinations	China, India, the former Soviet Union, Pakistan, all of the Middle East, Vietnam, most of Eastern Europe
Tier 4	Nations of concern	Iraq, Iran, Libya, North Korea, Cuba, Sudan, and Syria

Table 1. Country tier description

[1]This threshold will be raised to 28,000 MTOPS on February 26, 2001 and then 85,000 MTOPS on March 20, 2001.

[2]As of January 19, 2001, Tier 1 and Tier 2 were combined.

Tier Level	Jun 1991[A]	Feb 1994[A]	Jan 1996	Aug 1999	Feb 2000	Aug 2000	Jan 2001
Tier 1	195	1,500	No Limit	No Limit	No Limit	No Limit	No Limit
Tier 2	195	1,500	10,000	20,000	33,000	45,000	No Limit
Tier 3							
Civilian	195	1,500	7,000	12,300	20,000	28,000	85,000
Military	195	1,500	2,000	6,500[B]	12,500[C]	28,000[D]	85,000[E]
Tier 4	Virtual Embargo	Virtual Embargo	Virtual Embargo	Virtual Embargo	Virtual Embargo	Virtual Embargo	Virtual Embargo

Notes: A - Prior to Tier designations; B - Effective January 2000; C - Effective August 2000; D - Effective February 2001; E – Effective March 2001

Table 2. MTOPS license exception limits

B. Analysis of Computer Hardware Controls

The original intent of the MTOPS-based policy was to restrict exports of high-end computer systems, otherwise known as supercomputers, to countries of national security and proliferation concern. In the early 1990s, the MTOPS metric easily differentiated high-end systems from low-end commodity computers. For example, the MTOPS rating for a single-processor Cray C90 supercomputer was over fifty times that of a desktop Intel486 personal computer. As the decade progressed, rapid improvements in commodity, mass-market computer and communications technologies were to have a significant impact on high performance computing and the controls used to restrict computer exports. Consequently, MTOPS increases in commodity systems such as Intel-based personal computers outpaced Cray supercomputers (see Figure 1). Advances in parallel computing, both in systems software and hardware, enabled commodity computer manufacturers to build high performance computer systems from aggregations of the same commodity computer components (i.e., processors, system boards, controllers, hard disks, etc.) that are used in a wide variety of products, from desktops to servers to high performance computers. Today's top high performance computers contain thousands of commodity processors. These technology trends have removed the ability of the MTOPS metric to distinguish between the classical supercomputers and low- and mid-range servers. Controls that were intended to restrict the export of supercomputers are now ineffective because of the widespread availability of commodity servers having performance levels comparable to these expensive high-end systems.

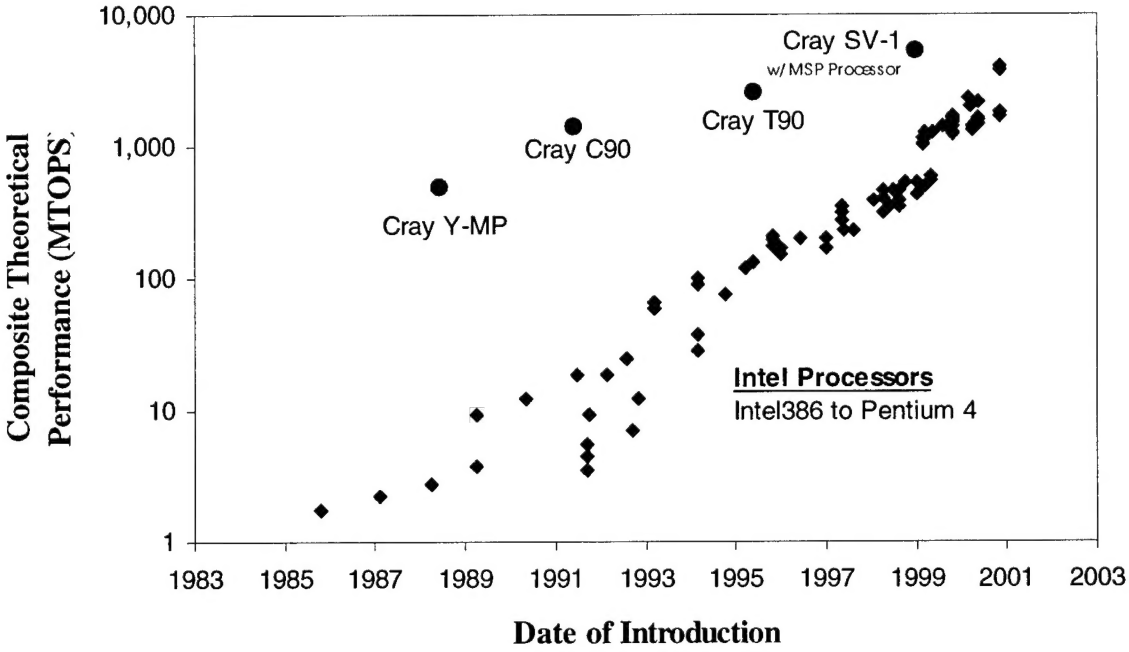

Figure 1. Trends for single-processor Cray supercomputers and Intel PCs[3]

From 1991 until the late 1990s, controls based upon MTOPS served as an effective means of restricting the export of high performance computers. During this period, high performance computers were sold as integrated products requiring manufacturer installation, maintenance, and support. Although large systems were built from commodity components, high performance computers were sold as a "single system image", requiring specialized hardware interconnections, chassis, and software. Vendors also supplied services such as software patches to repair system bugs, updated versions of the operating systems, hardware repair, and system upgrades to add processors, hard drives, and memory. User reliance on vendor installation and support meant that manufacturers could control distribution and therefore exports. Since users could not readily upgrade their systems, MTOPS thresholds served to limit the computational capability. Vendors, not eager to incur stiff penalties for unlicensed exports, typically adhered to export control regulations.[4] The government's ability to influence vendor delivery coupled with a distinguishing metric for computer performance were the two important factors that made controls effective and enforceable.

In 1994, Thomas Sterling and Don Becker at NASA initiated research to build high performance computers consisting of desktop computers and commodity network switches [4]. Although similar research has been conducted over the last two decades, the success of Sterling and Becker in building such systems, known as *Beowulf clusters,* inspired considerable research within the high performance computing community.

[3] Data sources: http://www.intel.com/pressroom/kits/processors/quickrefyr.htm, http://support.intel.com/support/processors/CTP.HTM, and Cray, Inc.

[4] On July 31, 1998, a major computer vendor was fined over $8.5 million for exporting high performance computers to a Russian nuclear weapons laboratory.

Cluster technology represents a particular challenge to effective export controls, since system software is freely available on the Internet and hardware components are readily available worldwide, in millions of units. This technology did not have a significant impact on export controls of computer hardware until the end of 1999, at which time cluster high performance computers matured into a commercially viable alternative to commodity systems. For the first time, in January 2001, export controls on computers were revised using clustered systems, rather than supercomputers or commodity multiprocessor systems, to determine the basic control threshold.

Over the last fifteen months, cluster solutions have been awarded in several open procurements for high performance computer systems. A prominent example occurred in November 1999, when the Forecast Systems Laboratory in Boulder, Colorado, awarded a small company, High Performance Technologies, Inc., an integration contract to provide a cluster system, which was rated at over 250,000 MTOPS.[5] (See Figure 2.) The implication of this and similar procurements was that any Tier 1, 2, or 3 country with the necessary technical expertise could build a high performance computer to an arbitrarily large MTOPS level.

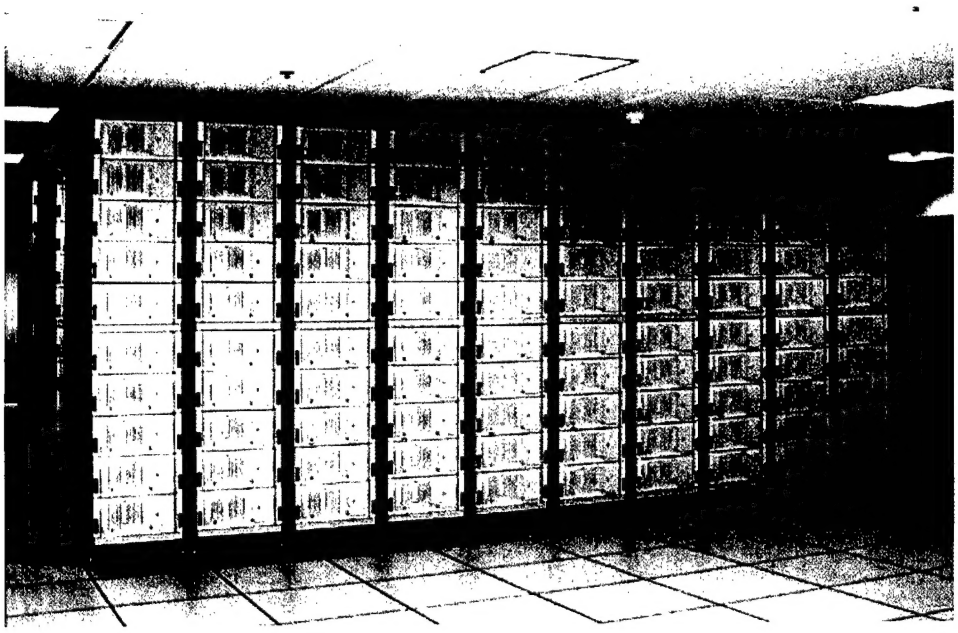

Figure 2. High performance PC cluster at Forecast Systems Laboratory, Boulder, Colorado. MTOPS rating over 250,000.

[5] The Forecast Systems Laboratory cluster is comprised of 277 Compaq XP1000 computers, each containing one 667-MHz Alpha EV67 processor, connected via Myricom Myrinet network switches. A presentation describing this system can be found at http://www-ad.fsl.noaa.gov/ac/Jet/usage000518/UsingJet.htm.

Starting around 1996, government agencies such as DOE, DoD, and NASA conducted aggressive programs to move critical applications software from Cray vector supercomputers onto integrated commodity multiprocessor systems (e.g., IBM SP2, Cray T3E, SGI Origin 2000). Today, many important applications run effectively on these systems, in particular, systems with distributed memory architecture – the same architecture used in many cluster systems. The transition of software from distributed memory multiprocessor systems to clusters is becoming routine. For example, the Forecast Systems Laboratory procurement used applications and systems software benchmarks to assess the performance of proposed hardware solutions. The results of these benchmarks, which included weather models, demonstrated the superior cost-performance of the awarded cluster over similar offerings by major computer manufacturers. Today, many applications codes can effectively utilize clusters.[6]

According to Goodman, et al. [5], an effective control policy for high performance computers requires the ability to identify certain characteristics of computers that permit effective forms of control. In the early 1990s, the MTOPS metric effectively captured such characteristics, since high performance computer systems were sold as integrated products and were difficult to assemble from components. Today, cluster technology, open source software, and improved component interoperability enable users to easily integrate commodity hardware into large high performance computer systems.[7] Thus export control limits, based on an overall system performance metric such as MTOPS, can be easily circumvented. The appendix at the end of this report describes several alternative approaches for controlling hardware. Unfortunately, given the state of technology, no set of key characteristics was identified that would permit effective control.

C. Recommendations for Computer Hardware Controls

The conclusion based on the previous analysis is that license exception limits based on MTOPS do not restrict foreign access to high performance computing. Controls, which are no longer effective, should be removed. The emphasis of government control policy should shift from hardware controls to protecting critical software applications, as discussed in the following section. Therefore, we recommend removing the license exception limits for Tier 3 exports. However, we also recommend retaining Tier 4 controls via the current virtual embargo and continuing controls on computer *technology* (e.g., lithography). Finally, we recommend retaining Enhanced Proliferation Control Initiative (EPCI) controls, which provides authority for the government to block exports of computers at any level in cases involving exports to or at risk of diversion to end-uses or end-users of proliferation concern (e.g., foreign nuclear weapons laboratories).[8]

---

[6] Note that we are not asserting that all codes will run effectively on clusters. Some software will demonstrate better cost-performance on vector systems, shared memory architectures, or more highly integrated distributed memory parallel systems.

[7] In October 1999, Cornell Theory Center built a cluster, consisting of 256 Pentium III processors, in less than one day.

[8] While EPCI restricts direct vendor support for WMD applications, determined organizations can circumvent this control through re-shipping. EPCI controls would only prevent the export of large HPC systems requiring direct manufacturer support for installation and maintenance. Considering the devastating impact of WMD, retaining EPCI measures is considered prudent.

II. Protection Measures for Applications Software

A. Background

Over the last twenty years, the United States has developed critical national security applications software to exploit the growing availability of high performance computers. This software provides a wide range of capabilities, including radar signature predictions of aircraft, computational fluid dynamics models for improved ship and submarine design, and image processing software for sensor systems. This software is typically government owned or controlled. Much of this software is the result of substantial investments in development, testing, and verification. Model validation is achieved through expensive experimentation, relying on the incorporation of certain parameter and data sets that are specific to weapons and warfighter systems. Thus, applications software is neither quickly nor easily duplicated by an adversary. Over the years, the government has taken measures to protect this software. Government regulations are in place that guide the policies used to ensure that national security applications software is released and distributed to the appropriate end-users. However, our efforts in protection need to be stronger given that we can no longer count on our adversaries being hindered by the lack of hardware computing power.

B. Analysis of Software Protection Technologies and Controls

In addition to classification guidelines, the International Traffic in Arms Regulations (ITAR) and the Export Administration Regulations (EAR) are the main regulations that guide release and distribution procedures for military and dual-use software. Regulations are quite clear regarding controls on software. For example, the export of application software specifically used to design a weapon system is strictly controlled, requiring an export license to all destinations. However, the regulations are complicated and laboratory and center security personnel are often unfamiliar with ITAR and EAR issues related to software. Clearly, better guidance and improved information dissemination would benefit the defense software development community.

Flexible tools that prevent unauthorized use of critical application codes are also needed. For example, we need protection technologies to ensure that selected software runs only on intended computers, and that the source and machine code cannot be modified or reverse engineered. Currently, the strongest measures involve distributions of versions of the executable code that are "locked" to particular machines by the inclusion of added hardware and software. Some protection methodologies are available from the commercial sector, but would not withstand determined efforts by trained adversaries. Defeat mechanisms for some protection technologies are already freely available on various "hacker" web sites. For dual-use codes, less intrusive protection technologies are required. Since the military applicability of dual-use software can vary greatly, users must be able to tailor protection techniques to the national security sensitivity of the software.

C. Recommendations for Software Applications

We recommend initiating a government activity to conduct research and development in software protection technology, provide related training, and support the insertion of protection technologies into critical national security applications software. In addition, this organization would verify and validate protection technologies, and facilitate the deployment of protection tools into the software development community. Information dissemination would also serve as an important component of this activity, informing the community about advances in software protection and increasing awareness of ITAR and EAR issues. To lay the groundwork for this activity, we have initiated a study to assess current approaches for controlling software, identify strengths and weaknesses of the current measures, and develop recommendations.

III. Summary

We have provided a rationale as to why export controls based upon the MTOPS metric are no longer effective. Our recommended hardware strategy is to abandon MTOPS controls for high performance computers. Our recommended software strategy involves initiating an R&D activity to develop, test, and verify technologies and methodologies that protect and limit end-use while minimizing the burden on an authorized end-user. Also, we propose improving the effectiveness of current policies used to guide the release and distribution of software through training and improved information dissemination. We believe that these recommendations will lead to an export control strategy for high performance computing that is effective, focuses our resources on the most important and urgent issues, and ensures that the United States maintains leadership in high performance computing.

IV. Acknowledgement

The authors wish to thank Dr. Alfred Brenner, Mr. Cray Henry, Dr. Eric Landree, Dr. Steven King, and Mr. William Gabor for their assistance in reviewing this document.

V. References

[1] Hambert, T., "Managing Busted Barriers - Export-Control Roadblock," Electronic Business Magazine, June 1999.

[2] Brenner, A.E., and Howes, N.R., "The Round Table on Computer Performance Metrics for Export Control: Discussion and Results," IDA Document D-2116, Institute for Defense Analysis Technical Report, December 1997.

[3] Export Administration Regulations, http://w3.access.gpo.gov/bxa/ear/ear_data.htm, Category 4, Pages 14-19.

[4] http://www.beowulf.org

[5] S. Goodman, P. Wolcott, and G. Burkhart, "Executive Briefing: An Examination of High-Performance Computing Export Control Policy in the 1990s," IEEE Computer Society Press, Los Alamitos, CA, 1996.

Appendix - Alternative Control Strategies for Computer Hardware

During this study, we examined the following options either as a replacement or augmentation to the current export control strategy. None of these options was determined to be an effective alternative.

A. Control by increasing MTOPS

The current approach is to coordinate increases in MTOPS license exception limits with advances in computer performance. At the start of this study, MTOPS limits were considered problematic. Frequent increases were required to stay ahead of commodity processors. For example, Intel processors were demonstrating annual MTOPS increases of over 50%. Also, Tier 3 civilian and military limits were being pushed upwards by 4- and 8-processor Intel servers. Low-end systems, not supercomputers, were driving control thresholds. Many within industry and the government were concerned about the deleterious effect of these controls on mass-market computer exports, especially to growing third world markets.[9] In addition, there was a perception that 2,000 MTOPS was a large amount of computational capability. Indeed, in 1992, two thousand MTOPS was considered a significant level of capability, costing over $2 million. Today, this level of capability can be purchased for less than $2,000. The original thrust of this study was to find a metric less sensitive to the geometric increases in processor performance. The technology available to cluster together low-MTOPS computers (i.e., below 6,500 MTOPS) into large high performance computer systems far exceeding license exception limits led us to the conclusion that any aggregated systems metric such as MTOPS could be easily circumvented, and was therefore an ineffective approach, and should be abandoned.

B. Control by number of processors

Several industry representatives advocated the replacement of the MTOPS metric with system processor count. Processor count is a simpler metric that, unlike MTOPS, is not dependent upon processor features and would not require frequent updates. Two important issues were identified: systems with less capable processors are penalized and the lack of a precise definition of a *processor*. For example, processors such as the IBM Power4, which contains two CPU functional units, can be counted as either one processor or two, depending on the definition. This approach was rejected since cluster technology enables one to easily circumvent export limits.

C. Control by software benchmarks

This proposal involved replacing the MTOPS metric with applications software benchmarks. Unlike MTOPS calculations, benchmarking is very labor intensive, costly to both implement and maintain, subject to interpretation, and susceptible to gaming. Due to cost and complexity, this approach was rejected.

---

[9] According to market data from International Data Corporation, Inc., the annual projected growth of computer sales in China and India is 22% and 16%, respectively.

D. Integrate key-encryption technology into computers to restrict end-use

This alternative strategy involved encouraging industry to integrate key-encryption into the important computer components (i.e., processors, controllers, and other circuitry) to restrict end-use. This concept would enable a vendor to determine the allowable end-use of a given system. For example, a system might be allowed to run Oracle at full system performance but unable to run FORTRAN codes. Successful implementation would require significant R&D and industry buy-in. This option was deemed too ambitious given its technical complexity and significant cost.

E. Control by bandwidth

A metric based on the peak single-processor-to-main-memory bandwidth was proposed as either an augmentation or replacement to MTOPS. The purpose of this metric was to control the export of traditional vector supercomputers such as the Cray T90. In addition, we considered a license exception limit of 20 Gigabytes per second. As demonstrated in Figure 3, this metric can be used to differentiate commodity-based systems such as the Compaq ES-40 from high-end vector systems such as the NEC SX-5. Since Cray no longer manufactures the T-90 and their next generation product, the SV-2, is under development, this control measure would currently restrict only the export of high-end Japanese systems. Obtaining international agreement for this metric was considered problematic, since Japanese exports would be impacted and United States exports would not.

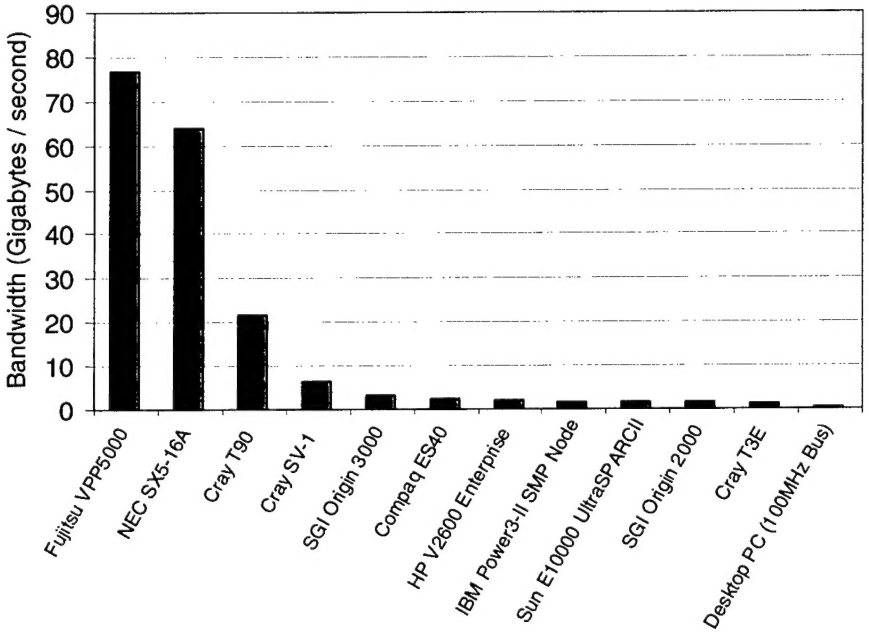

Figure 3. Peak Single-Processor-to-Main-Memory Bandwidth for Various Computer Systems